Design Principles for Ontological Support of Bayesian Evidence Management

Michael N. HUHNS, Marco G. VALTORTA, and Jingsong WANG
University of South Carolina, Columbia, SC, USA

Abstract. This chapter describes work on an integrated system that can assist analysts in exploring hypotheses using Bayesian analysis of evidence from a variety of sources. The hypothesis exploration is aided by an ontology that represents domain knowledge, events, and causality for Bayesian reasoning, as well as models of information sources for evidential reasoning. We are validating the approach via a tool, Magellan, that uses both Bayesian models and logical models for an analyst's prior knowledge about how evidence can be used to evaluate hypotheses. The ontology makes it possible and practical for complex situations of interest to be modeled and then analyzed formally.

Keywords. causal ontology, Bayesian reasoning, evidence management

1. Introduction

Much of the extensive work on ontologies to date has focused on modeling and representing the world of objects. The ontologies needed for our research supporting the management of hypotheses and evidence for analysts, however, must additionally model events and causality. Less work has been done on this aspect of ontologies. In this paper we show how concepts from a causal ontology can be used directly as variables in Bayesian networks and how the attributes of the causal concepts can be used in matching evidence to the variables. Moreover, subclass relationships in the ontology enable the extension of Bayesian reasoning over types.

2. Bayesian Reasoning for Evidence Management

There are numerous real-world situations about which an analyst might wish to hypothesize and investigate, but it would be impractical to encode all of them explicitly in a support system for analysts. Instead, our approach is to represent fragments of situations and provide a mechanism for combining them into a wide variety of more complete ones [1,2]. The combination occurs dynamically as evidence about a situation becomes available or as an analyst revises or enters new hypotheses. A situation fragment is represented as a Bayesian network with nodes for hypotheses, events, and evidence, and links for relating them. Our ability to combine the fragments into more complete situation models is dependent on having a consistent terminology in which the fragments are described. The focus of our work has been on (1) defining and representing the terminology, including terms of a domain and terms for evidence in that domain, (2) capturing new fragments from a variety of sources, and (3) incorporating the terminology and BN fragments into an integrated end-to-end tool, Magellan.

2.1. Recognizing and Representing Situations

Our objective is to be able to model and reason probabilistically about a wide variety of situations that might be of interest to analysts. Unfortunately, there are too many situations for system developers to encode *a priori*, which even if possible would make the resultant system too complex for analysts to use, and it is unrealistic to expect most analysts to be able to use the requisite formal mechanism to encode situations *a posteriori*. Instead, our approach is to represent small, common aspects of situations generically, and then provide a means to combine them dynamically into representations for real-world situations. We term the small generic situation aspect a *fragment*, and choose a first-order representation for it.

An example situation aspect that we might represent as a fragment would be a "suspicious transfer of money," with variables corresponding to banks, organizations, deposits, withdrawals, and the transferring agent. The fragment would be instantiated when evidence matched the variables, e.g., "a church attended by Syrians in Detroit deposited funds into a Michigan bank and the funds were transferred to a bank in Cairo." More precisely, each variable (node) in a fragment has a set of identifying attributes and their collective instantiated values specify a particular instance of a random variable. Because the evidence might be uncertain, there would be probabilities associated with the instantiated fragment, and we would treat the instantiated fragment as a Bayesian network. This is shown in Figure 1. Note that the probability distribution described in the Bayesian network is a joint distribution on the nodes only, not on the nodes and the attributes.

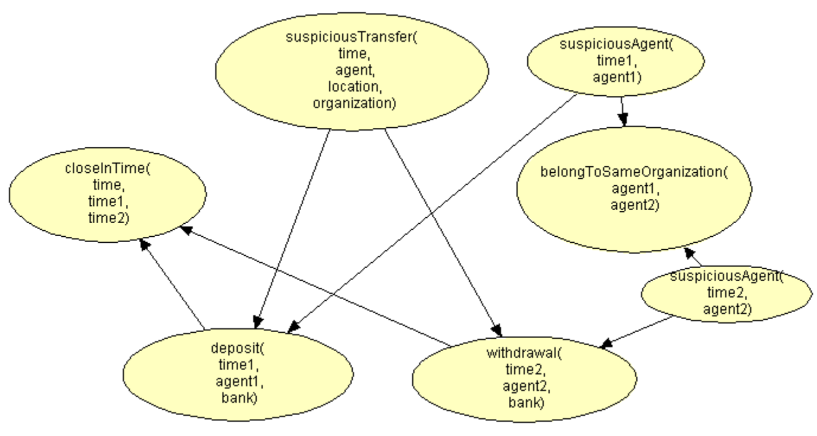

Figure 1. A commonly occurring part of a situation for a suspicious bank transfer of money, represented as an uninstantiated Bayesian network. Notice that the nodes (variables) have attributes, making them equivalent to concepts or classes in an ontology. If one or more items of evidence matched the nodes, then details of the evidence would be used to instantiate the attributes of the variables

An advantage of using fragments of situations instead of more complete situations is that many more situations can be represented efficiently. More precisely, *N* fragments can potentially be combined in *N!* ways to represent *N!* situations. The combining is guided by available evidence. For example, three other situations that we might represent as fragments are "purchases of weapons," "influencing an election," and "bribing a politician." If evidence matched one of these, and the resulting instantiated fragment had one or more variables in common with the money fragment, then we would merge the fragments at the point of the common variables to produce a representation of a more complete situation, such as "transferring money to influence an election." Note that fragments can be merged only if the attributes of their common variables unify. Also note that it is not necessary for the fragments to have any variables in common in order to merge them and represent larger situations. As a result, the fragments could represent situations such as "bribing a politician to influence an election" and "purchasing weapons to influence an election." Further, because each fragment could be instantiated multiple times, we could represent several different money transfers being used to purchase weapons. Our system, Magellan, considers all of the possible situations that are consistent with available evidence. Magellan then performs Bayesian reasoning on whichever complex situation representation resulted from instantiating fragments with the available evidence and integrating those fragments. The overall process for merging instantiated fragments and reasoning over them is shown in Figure 2.

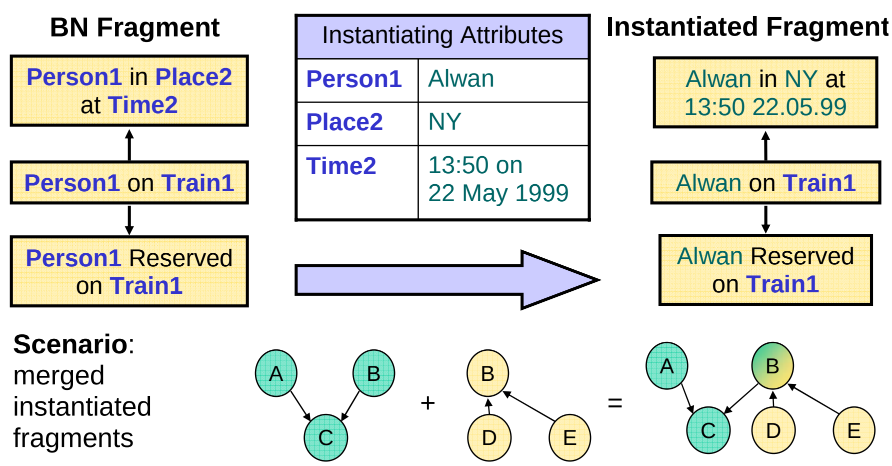

Figure 2. Fragments (left side)—in this case about people being on a train, having made reservations on a train, and being at the same location as the train—are merged based on the evidence (center) that instantiates them

2.2. Capturing the Terminology and Prior Knowledge for a New Domain

A key activity of an intelligence analyst is to distinguish among competing hypotheses, determine the likelihood of their occurrence, and reduce the uncertainty in the outcomes of the hypotheses, upon which decision makers will then base their decisions.

Hypothesis outcomes[1] are related to observable evidence via direct or indirect causal relations, and therefore ontological support for analysts should involve cause-and-effect. This is best supported by an ontology emphasizing events and their causal relationships, along with a hypothetical world of possible events, actions, and causes. However, causal relationships must be interpreted in the context of the state of the real world—primarily consisting of objects and their physical properties—which can be represented in a conventional ontology, such as those that are part of SUMO. The evidence for reasoning about hypotheses can come from a variety of sources, and the acquisition of evidence and events from these sources must also be represented, constituting a third kind of ontological representation describing the information sources. Figure 3 depicts the three ontological models we use for (1) modeling situations and relating them to (2) background knowledge about the state of the world, and (3) acquiring evidence, all of which enables an assessment of the likelihood of the situations using Bayesian reasoning.

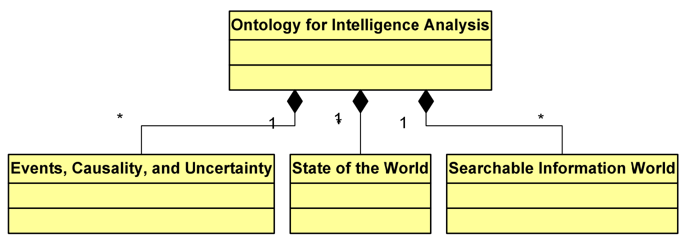

Figure 3. An ontology for intelligence analysts has three related parts, corresponding to (1) the world of causality and hypothetical events needed for Bayesian reasoning, (2) the real world of things needed to model situations, and (3) the world of information and information sources needed for evidence management

A situation might represent an analyst's query or, more generally, provide context and support for a hypothesis. A situation would be comprised of one or more items of interest and each such item of interest has information provided by several information sources. An item of interest may be specialized to Person, Organization, Event, or Place, and of particular interest would be items relating events involving people at significant places. Information sources can be maps, images, reports video, audio, email, websites, and database records. Typically, an item of interest would have many information sources describing aspects of that item, for example a meeting held by members of a suspected terrorist organization might be described by audio, video, and email surveillance or reports by insiders. Our tool, Magellan, uses Protégé [3] (see Figure 7) for capturing the ontologies, RDF (Resource Description Framework [4]) for representing the terminology, XMLBIF (eXtensible Markup Language Bayesian Interchange Format [5]) for representing the causal relationships, and RDF and SPARQL (a query language for RDF [6]) for requesting evidence from information sources. It also makes use of logical, non-probabilistic models, as shown in Figure 4 and described next.

[1] In our ontology, an outcome is thus an important and necessary property ("slot" in Protégé) for hypotheses and, indeed, for any concept that may be in a causal relationship. The relationship is a link in a Bayesian network.

2.3. Situation Fragments Represented by Logical Models

Our objective is to produce models of systems and situations that will be sufficiently accurate that they can be used—where appropriate—to predict future states, to understand operations, to illuminate the factors relevant to decisions, and to control behaviors. We have realized that some knowledge is more easily and naturally represented in the form of statements in a logic language and some is more naturally represented in a Bayesian-network formalism. For example, logic is best for expressing

- Class-subclass statements, such as "C4 is an explosive"
- Part-whole statements, such as "triggers are part of IEDs"
- Definitional statements, such as "triangles have three sides"
- Temporal statements, such as "3:00 p.m. occurs before 4:00 p.m."
- Spatial statements, such as "Irbil is located in Kurdistan"

Other knowledge is probabilistic, such as

- "Terrorist cell X planned the bombing"
- "Suspect Y met with cell leader Z in Syria last March"

Our resultant reasoner takes advantage of the strengths of each formalism, while integrating them into a single coherent system.

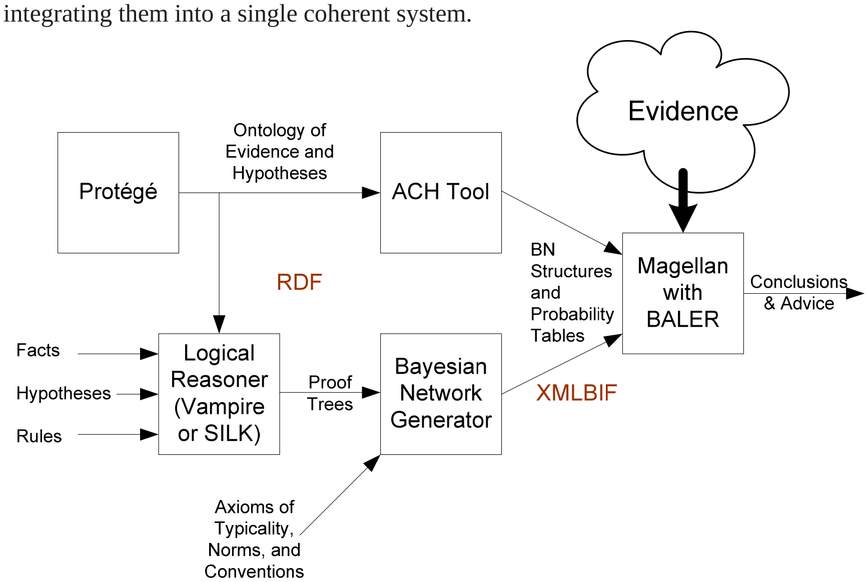

Figure 4. The BALER framework for integrating logical models with probabilistic models, with an ontology developed in Protégé providing a consistent vocabulary for all domain concepts

An example of the situations that can be represented by such an integrated system is shown in Figure 5. This system would help analysts confront problems of *credibility*, *relevance*, *contradictory evidence*, and *pervasive uncertainty*, using

- A unique combination of the power of logical and probabilistic reasoning
- Numerical analysis of competing hypotheses
- Automated linking of relevant evidence

- Automated propagation of uncertainty values: good arguments from uncertain data still add strength to a conclusion
- Robust reasoning over contradictory information allows analysts to exploit maximal amounts of information
- A provision for analysts to enter their own knowledge directly, allowing the system to learn from its users
- The use of probabilities to quantify belief in hypotheses to support optimal decision making according to the principle of maximum expected utility.

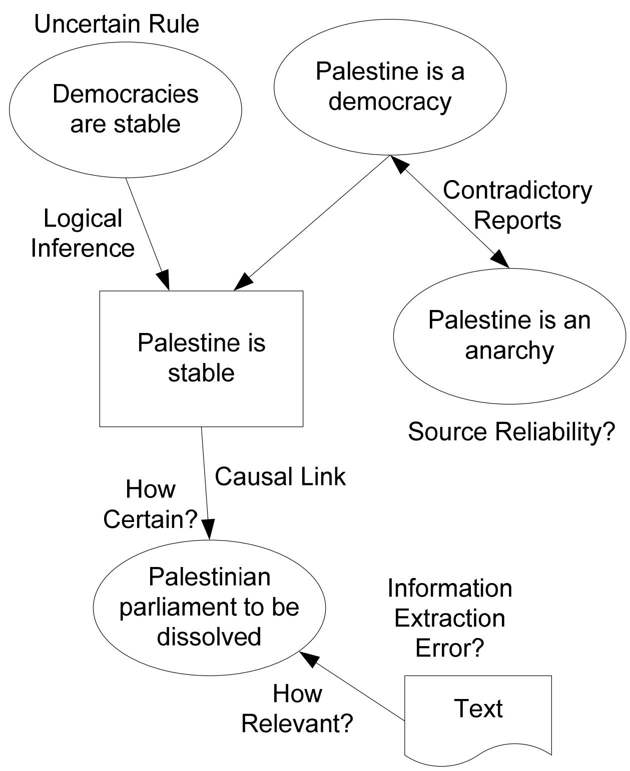

Figure 5. An example illustrating the need for both Bayesian and logical reasoning

Formal logical tools are able to provide some amount of reasoning support for information analysis, but are unable to represent uncertainty. Bayesian network tools represent probabilistic and causal information, but in the worst case they scale as poorly as some formal logical systems and require specialized expertise to use effectively [7]. The framework (BALER) we have developed for intelligence reasoning incorporates the advantages of both Bayesian and logical systems [8]. The framework includes a formal mechanism for the conversion of automatically generated natural deduction proof trees into Bayesian networks. This is indicated by the information flow shown in Figure 6. We have proven that the merging of such networks with domain-specific causal models forms a consistent Bayesian network with correct values for the formulas derived in the proof. In particular, we show that hard evidential updates (see Section 2.5) in which the premises of a proof are found to be true force the conclusions of the proof to be true with probability one, regardless of any dependencies and prior probability values assumed for the causal model.

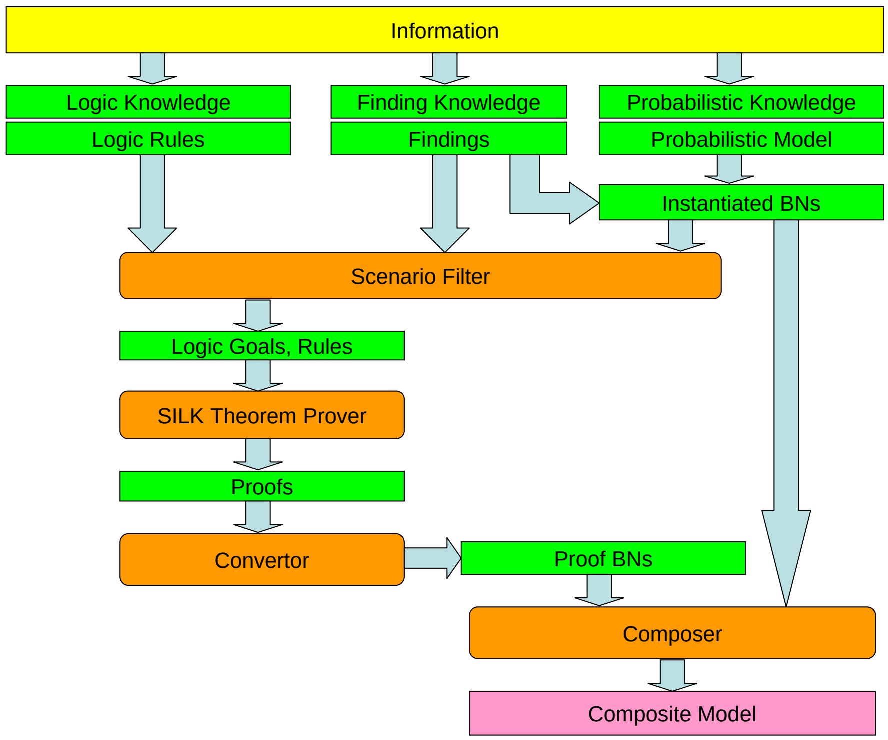

Figure 6. The BALER software process flow, which is supported by the tripartite ontology of real world concepts, events, and information sources

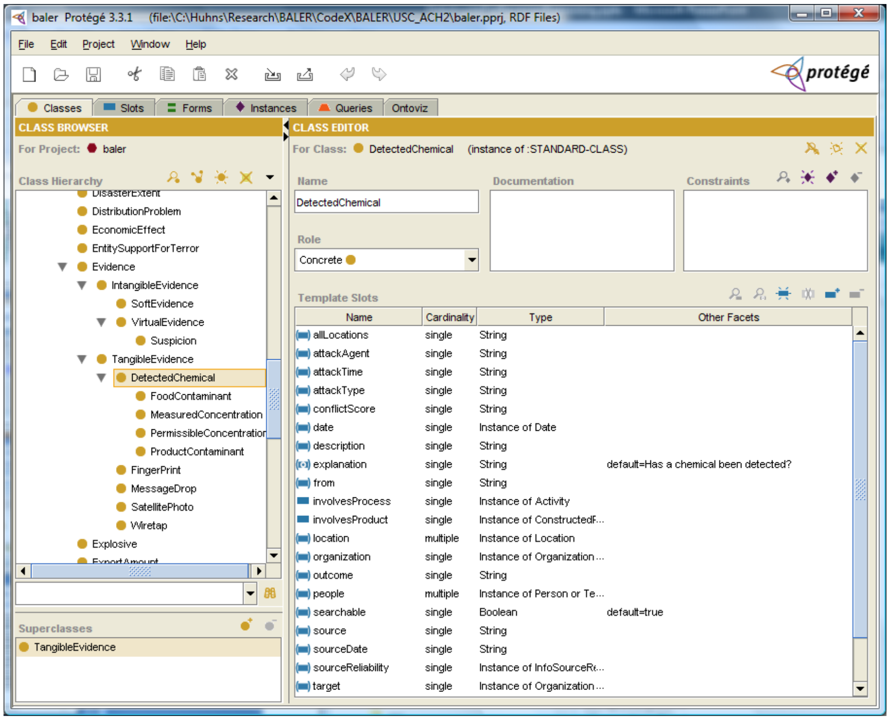

Figure 7. Protégé is used to enter the ontology concepts that form the basis for representing situations and evidence

2.4. Causality

Causality is a special relationship among events for which certain properties hold probabilistically. For example, causality is logically irreflexive and asymmetric, but probabilistically transitive. Causality, like the relation *subevents*, generates a strict partial order among events. Causal models are very useful, because they allow prediction of the effect of interventions [9,10]. Our interest is in a causal Bayesian network.

A *causal Bayesian network* consists of a causal graph, a directed acyclic graph (DAG) expressing causal relationships, and a probability distribution respecting the independence relation encoded by the graph [8]. A link between two nodes in a Bayesian network is often interpreted as a causal link. However, this is not necessarily the case. When each link in a Bayesian network is causal, then the Bayesian network is called a causal Bayesian network or Markovian model. A Markovian model is a popular graphical model for encoding distributional and causal relationships. To summarize, a Markovian model consists of a DAG G over a set of variables $V = \{V_1; \ldots ; V_n\}$, called a causal graph and a probability distribution over V that has some constraints on it. The interpretation of such a model consists of two parts: the association of the variables to events and the assignment of probability distributions to the links. For causality, variable assignment must satisfy the obvious constraint that

(Event A *causes* Event B) ➔ ($time_A$ < $time_B$)

The probability distributions must satisfy two constraints. The first constraint is that each variable in the graph is independent of all its non-descendants given its direct parents. The second constraint is that the directed edges in *G* represent causal influences between the corresponding variables. A Markovian model for which only the first constraint holds is called a Bayesian network, and its DAG is called a Bayesian network structure. This explains why Markovian models are also called causal Bayesian networks. As far as the second condition is concerned, causality requires that, when a variable is set, the parents of that variable be disconnected from it: this is called the excision model of causality.

In our prototype tool, Magellan, new variables are added to the causal and event portion of an analyst's ontology using Protégé, so that all of the nodes in a Bayesian network fragment are represented in a standard and consistent terminology. We extend SUMO with this terminology, so that we can take advantage of SUMO's existing description of general knowledge of the world. Each variable has a set of identifying attributes, which are used to combine fragments (fragments can be combined only if their attributes unify) [1,2].

Probabilities are assigned to events in the fragment by performing experiments, estimating beliefs, or counting outcomes. Once assigned, they are updated by conditioning on evidence using Bayes rule and the laws of probability. The fragments are stored in a repository, where they can be matched with evidence and combined with other fragments to produce models of situations that are as complete, accurate, and specific as possible.

2.5. Evidence

Fragments are instantiated by evidence, which we define informally as information (perhaps wrong, perhaps incomplete) about what happened (events). For example, a bank clerk might be uncertain whether a money transfer was to a Cairo bank or a Boston bank. We represent in the information source ontology the level of credibility of items of evidence, and provide a Bayesian interpretation of credibility. Formally, we define *evidence* to be a collection of findings, each of which describes the state of a Bayesian network variable, and distinguish three kinds [7]:

1. A *hard finding* specifies that the variable has a particular value. For example, whether or not a money transfer occurred or whether or not a suspect is a terrorist

   (Male_TerroristSuspect = true)

2. A *soft finding* is a distribution on the states of a variable, usually corresponding to an "objective" statistical distribution that is not expected to change within a scenario [11]. For example, there might be an observation that 95% of terrorists are male (and 5% are not), i.e.,

   Q(Male_TerroristSuspect)=(0.95, 0.05)

3. A *virtual finding* is a likelihood ratio corresponding to the credibility associated to an evidence source, such as a witness. For example, witness Bill might have

observed a suspect entering a men's-only area of a mosque, which would be interpreted as 4-to-1 that the suspect is a male

```
L(Male_TerroristSuspect)=(0.8, 0.2)
```

Unlike soft findings, virtual findings allow for an update of the posterior probability of the evidence variable.

The relationships among the evidence types are shown in Figure 8.

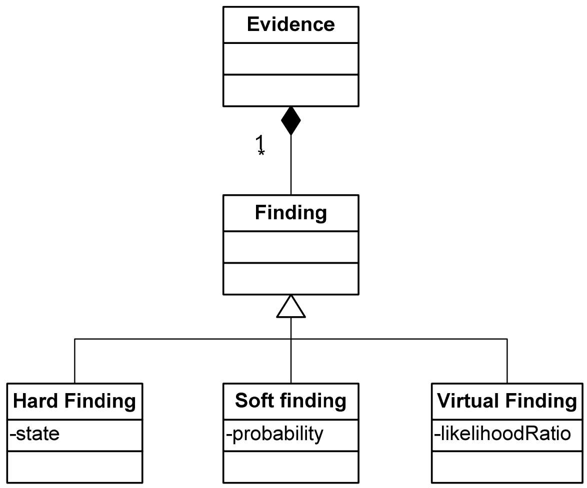

Figure 8. Evidence consists of a set of findings, which can be of three different types, hard, soft, and virtual

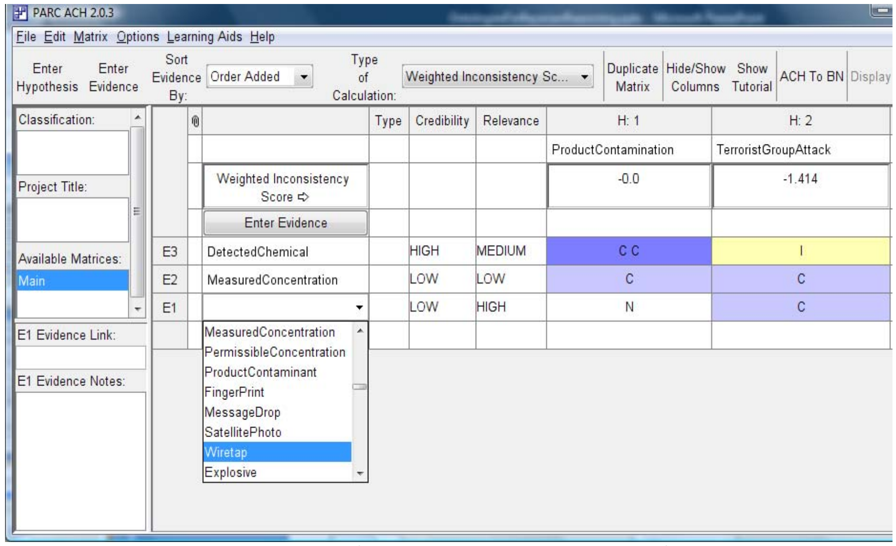

Figure 9. Magellan's extended ACH interface is integrated with the ontology of events through pull-down menus, i.e., each hypothesis (such as "TerroristGroupAttack") and each type of relevant evidence (such as "DetectedChemical") is a concept from the domain ontology

Our modified version of the tool ACH2 [12] is used by an analyst to enter the appropriate hypotheses and any initial evidence that might be available. The terminology available to the analyst is provided via drop-down menus as shown in Figure 9, where the menu entries are the ontology terms from our ontology developed in Protégé [3]. The use of terms from an ontology is essential for (1) enabling logical proofs to be constructed out of both new knowledge and prior knowledge, (2) taking advantage of known generalizations and specializations for reasoning and fragment matching, (3) guiding analysts in the kinds of concepts that can be used to represent hypotheses and evidence, and (4) enabling new fragments to be composed with existing fragments to represent situations more comprehensively.

The resultant Analysis of Competing Hypotheses (ACH) [13] matrix is converted automatically into a bipartite Bayesian network, with initial probabilities assigned based on the relevance factors assigned to cells of the matrix. An example of the network is shown in Figure 10. The network is saved into a repository of fragments, from where it can be retrieved for matching to evidence and then composed with other fragments.

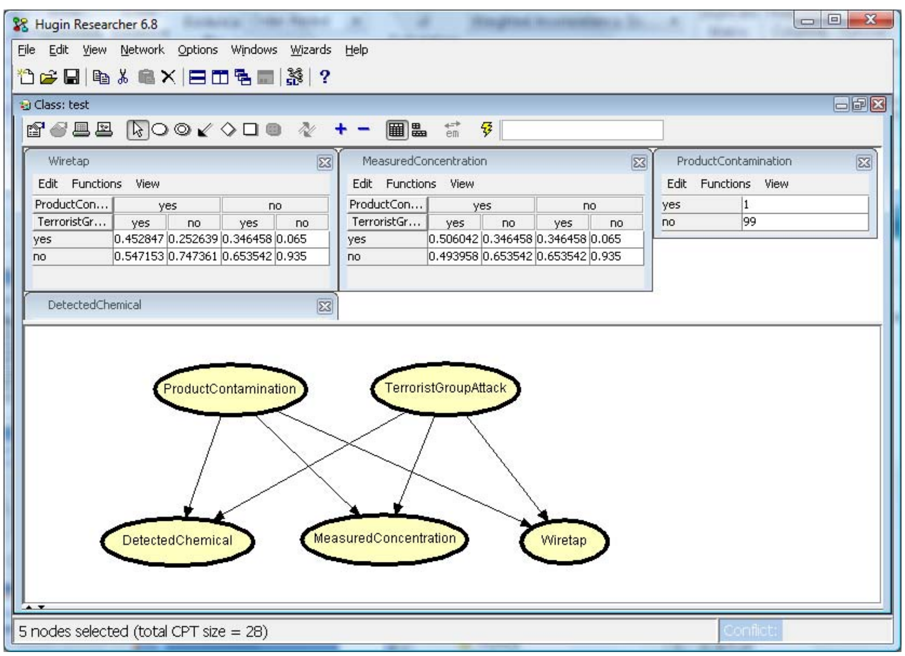

Figure 10. A Bayesian network fragment constructed automatically from an ACH matrix. The conditional probabilities needed for Bayesian reasoning are derived from the user-entered values in the matrix indicating whether or not a finding is consistent with an analyst's hypothesis

3. Use of Tripartite Ontology for Intelligence Analysis

Figure 11 shows an end-to-end architecture for Bayesian reasoning, which would be used as follows. The process might be triggered by the arrival of evidence in the form of a message, such as the following:

FBI Report Date: 10 April 2003. FBI: Abdul Ramazi is the owner of the Select Gourmet Foods shop in Springfield Mall. Springfield, VA. (Phone number 703-659-2317). First Union National Bank lists Select Gourmet Foods as holding account number 1070173749003. Six checks totaling $35,000 have been deposited in this account in the past four months and are recorded as having been drawn on accounts at the Pyramid Bank of Cairo, Egypt and the Central Bank of Dubai, United Arab Emirates. Both of these banks have just been listed as possible conduits in money laundering schemes.

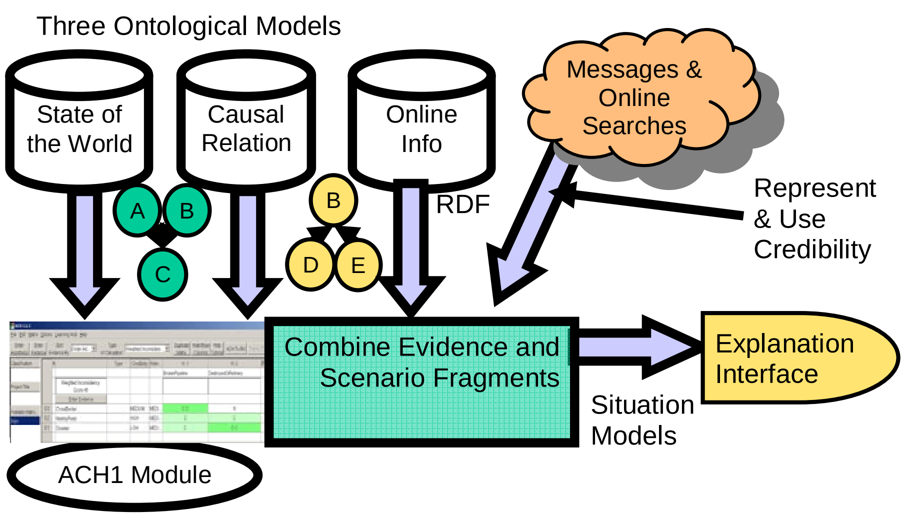

Figure 11. Magellan architecture for Bayesian Reasoning used to explore an analyst's hypotheses, indicating how the ontology makes it possible for evidence to be combined with generic situation fragments to produce models that can be reasoned over probabilistically to explain the evidence

Based on such a message, or based on a hypothesized situation that an analyst would like to investigate, an appropriate scenario represented as a Bayesian model is chosen by the analyst and a corresponding form is displayed listing initial evidence and the domain variables for the scenario. The evidence values for the variables can be supplied automatically from the triggering messages, by matching message terms with ontology concepts as shown in Figure 12, or can be entered by the analyst. Because the probabilities of the variables represented in a situation are updated to be consistent with the evidence at hand, the situation tracks the variables of interest to an analyst. When the probability of a particular value of a variable of interest becomes sufficiently high, an alert could be issued to the analyst.

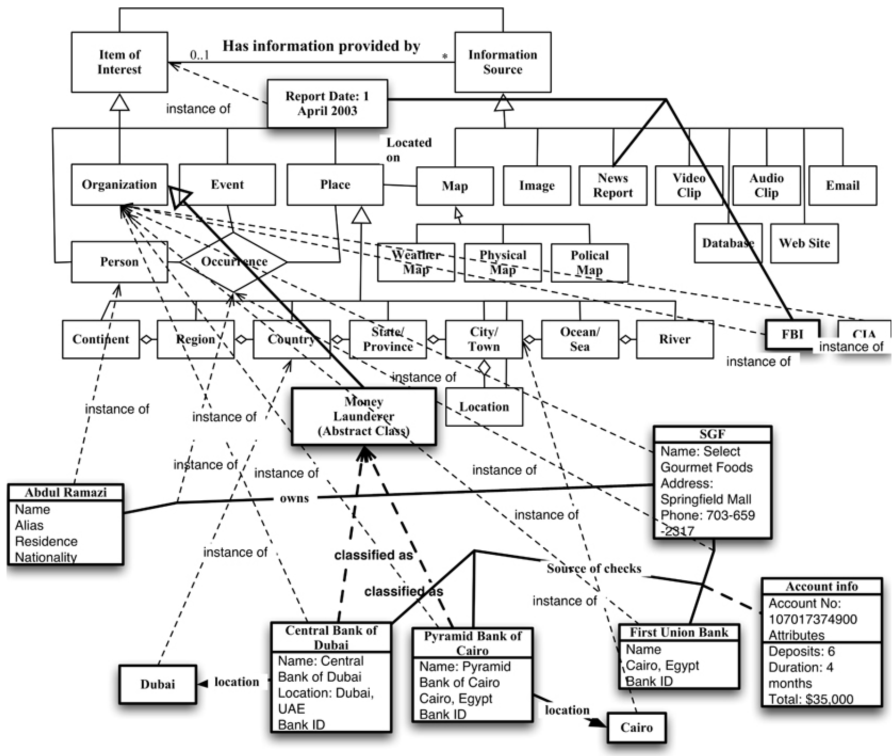

Figure 12. A small portion of the tripartite ontology indicating how an item of evidence would be classified and used to instantiate one or more fragments. Shown here is the outcome of representing the FBI evidence message above using the ontology that we defined in Protégé

Then the Bayesian reasoning component, using a value-of-information calculation, identifies the variables that have the most potential impact on the probability profile of a variable of interest. (Algorithm 1 contains the algorithm that we use to calculate the value of information for a chosen variable.) That is, it determines which pieces of evidence would be most useful in confirming or denying the analyst's hypothesis. Such especially informative variables can then become the subject of focused queries. A request for this evidence is sent to the analyst, who returns the result to the Bayesian reasoner for incorporation into the situation, and the likelihood of the analyst's hypothesis is reassessed. The process is repeated until the analyst decides to stop or there is no more evidence available that changes the plausible outcomes.

Algorithm 1. Value-of-Information Calculation

- Let *V* be a variable whose value affects the actions to be taken by an analyst. For example, *V* indicates whether a bomb is placed on a particular airliner.
- Let *p(v)* be the probability that variable *V* has value *v*.
- The entropy of *V* is:

$$\mathrm{H}(V) = -\sum_{v \in V} p(V = v) \log\left(p(V = v)\right)$$

- Let *T* be a variable whose value we may acquire (by expending resources). For example, *T* indicates whether a passenger is a known terrorist.
- The entropy of *V* given that *T* has value *t* is:

$$\mathrm{H}(V|t) = -\sum_{v \in V} p(V = v|T = t) \log\left(p(V = v|T = t)\right)$$

- The expected entropy of *V* given *T* is:

$$\mathrm{E}[\mathrm{H}(V|t)] = \sum_{t \in T} p(T = t)\, \mathrm{H}(V|t))$$

- The value of information is then:

$$VOI(V) = -(\mathrm{E}[\mathrm{H}(V|t)] - \mathrm{H}(V))$$

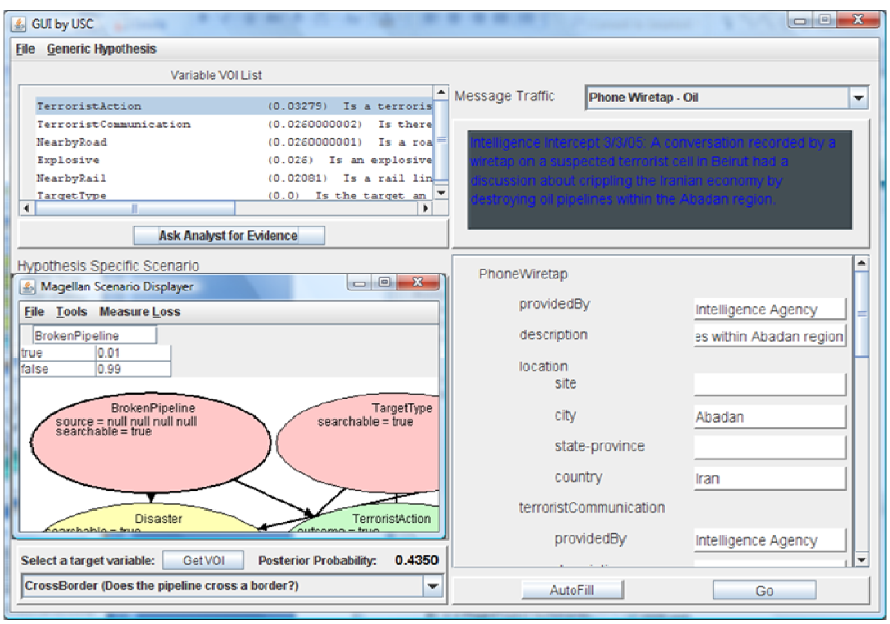

Figure 13. The Magellan interface showing an evidence message (upper right), the ontology concepts it contains (lower right), the fragments that it instantiates, composed into a situation (lower left), and the posterior probability for an hypothesis about the situation

4. Evaluation

An early anecdotal evaluation of Magellan was conducted at NIST. The evaluators (three naval reservists with a background in intelligence analysis) tested the hypothesis generation aspect of the system for four hours. In this test, the analysts were presented with several items of evidence (similar to the FBI Report in of section 3) and asked to generate hypotheses, using an interface such as is shown in Figure 13. After they had finished, they were shown hypotheses generated by Magellan and were asked to rate these hypotheses in comparison to the ones they had generated. The NIST summary of the evaluation indicated that the analysts generated more hypotheses than Magellan and that Magellan's hypotheses did not take into account all the possible variables. However, analysts' ratings for Magellan-generated hypotheses are equal to the ratings for the analyst-generated hypotheses in 1/3 of the cases. In 7/9 cases the ratings for the Magellan-generated hypotheses were given mid-level ratings or higher.

5. Discussion

The key features of our approach to reasoning about evidence are the ability to model fragments of abstract situations, to base the models on concepts from a causal ontology, to use a combination of both logic and probability for reasoning about the models, to ground situation models by instantiating the ontological concepts in the fragments with evidence, to compose the instantiations of situation fragments into complete situation models based on evidence, and to analyze the resultant situation models for sensitivity and surprise. The heart of our approach is Bayesian reasoning. However, there are alternative approaches for reasoning over uncertain evidence about ontological concepts, notably Pronto [14,15], a probabilistic extension to OWL [16], and P-Classic [17].

Pronto provides reasoning services for knowledge bases containing uncertain knowledge. It extends the Pellet reasoner by enabling probabilistic knowledge representation and reasoning in OWL ontologies. Pronto represents uncertainty by probability intervals, instead of point probabilities and tables of conditional probabilities, as in Bayesian networks. The advantages of the Bayesian approach are

- Bayesian networks directly support causality, which to do the equivalent in Pronto would require an additional logical theory.
- Both approaches can handle logical conflicts, but Pronto relies on a mechanism of model ordering via the use of preferences, whereas Bayesian networks make use of explicit models that describe the conflicts, so that they can be reasoned about in the same way as non-conflicting evidence.
- As evidence about an uncertain variable accumulates, the variable's probabilistic interval becomes wider and it becomes more difficult to base a decision on the variable.
- Probabilistic interval updating as done in Pronto is more complex than the updating of point probabilities in Bayesian reasoning.

In P-Classic, which supports conditional probabilities as in Bayesian networks, links represent subclasses, as opposed to representing causality. P-Classic is most useful for problems of identification, i.e., given some uncertain features about an unknown concept *x*, it can conclude that *x* is most likely an instance of class **Y**. The work on probabilistic extensions to OWL by Ding and Peng [16] improves on P-Classic by focusing on *formal* rules for translating OWL ontologies into Bayesian networks. Note that our ontology is not itself probabilistic and we do *not* translate it into a Bayesian network—we just use concepts from it in Bayesian networks and ensure that the Bayesian networks are consistent with the causality knowledge in the ontology.

6. Conclusion

Our work is predicated on the observation that ontologies make it easier for tools to interoperate. We have found that our ontologies need to describe both the physical world and the on-line information world, because our reasoning system relies on the relationships and links between both kinds of domains. The reasoner, BALER, enables first-order logic sentences to be combined with Bayesian networks by generating Bayesian networks for any first-order natural deduction proof (that uses the Reeves-Clarke inference rules). This exploits the complementary powers of both logical and Bayesian representations.

7. Acknowledgements

This work was funded in part by the Disruptive Technology Office Collaboration and Analyst System Effectiveness (CASE) Program, contract FA8750-06-C-0194 issued by Air Force Research Laboratory (AFRL). The views and conclusions are those of the authors, not of the US Government or its agencies.

References

[1] John Cheng, Ray Emami, Larry Kerschberg, Eugene Santos, Jr., Qunhua Zhao, Hien Nguyen, Hua Wang, Michael Huhns, Marco Valtorta, Jiangbo Dang, Hrishikesh Goradia, Jingshan Huang, and Sharon Xi, "OmniSeer: A Cognitive Framework for User Modeling, Reuse of Prior and Tacit Knowledge, and Collaborative Knowledge Services," *Proceedings of the 38th Hawaii International Conference on System Sciences*, HICSS38, 2005.

[2] Katherine Laskey and Suzanne Mahoney, "Network Fragments: Representing Knowledge for Constructing Probabilistic Models," in *Proceeding of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, AAAI Press, 1997, 334-341.

[3] The Protégé Ontology Editor and Knowledge Acquisition System, http://protege.stanford.edu/.

[4] Resource Description Framework (RDF), http://www.w3.org/RDF/.

[5] Fabio G. Cozman, "The Interchange Format for Bayesian Networks," 1998. http://www.cs.cmu.edu/~fgcozman/Research/InterchangeFormat/

[6] SPARQL Query Language for RDF, http://www.w3.org/TR/rdf-sparql-query/.

[7] Marco Valtorta and Yimin Huang, "Identifiability in Causal Bayesian Networks: A Gentle Introduction," *Cybernetics and Systems* **39:4** (2008) 425-442.

[8] Marco Valtorta, John Byrnes, and Michael Huhns, "Logical and Probabilistic Reasoning to Support Information Analysis in Uncertain Domains," *Proceedings of the Third Workshop on Combining Probabilty and Logic (Progic-07)*, Canterbury, England, September, 2007, 5-7

[9] Yimin Huang and Marco Valtorta, "Pearl's Calculus of Intervention is Complete," *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI-06)*, 2006, 217-224.

[10] Judea Pearl, *Causality: Modeling, Reasoning, and Inference*, Cambridge, England: Cambridge University Press, 2000.

[11] Marco G. Valtorta, Y.-G. Kim, and Jirka Vomlel, "Soft Evidential Update for Multiagent Systems," *International Journal of Approximate Reasoning* **29:1** (2002) 71-106.
[12] Peter Pirolli and Lance Good, "Evaluation of a Computer Support Tool for Analysis of Competing Hypotheses," UIR Technical Report, Palo Alto Research Center, 2004.
[13] Richards J. Heuer, Jr., *Psychology of Intelligence Analysis*, Center for the Study of Intelligence (at http://www.cia.gov/csi/books/19104/index.html), 1999.
[14] Thomas Lukasiewicz, "Probabilistic Description Logics for the Semantic Web," INFSYS Research Report 1843-06-05, Institut Für Informationssysteme, Technische Universität Wien, March 2007.
[15] Pronto--A Probabilistic Reasoner for OWL DL and Pellet, http://pellet.owldl.com/pronto.
[16] Zhongli Ding and Yun Peng, "A Probabilistic Extension to Ontology Language OWL," *Proc. 37th Hawaii International Conference on System Sciences*, HICSS37, 2004.
[17] Daphne Koller, A. Levy, and Avri Pfeffer, "P-CLASSIC: A Tractable Probabilistic Description Logic," in *Proc. AAAI-97*, AAAI Press, 1997, 390-397.